

Threat Paper

Federal Computer Incident Response Center (FedCIRC)

**Defense Tactics for Distributed
Denial of Service Attacks**



This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

**Federal Computer Incident Response Center
7th and "D" Streets S.W.
Room 5060
Washington, DC 20407**

This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

Acknowledgements

FedCIRC wishes to thank Global Integrity, Cisco, CheckPoint, Dr. Fred Cohen, David Dittrich, and Craig Huegen for their assistance in the preparation of this paper. Additionally, multiple members of Global Integrity Corporation contributed to this paper during formal and informal discussions.

This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

Disclaimer

This paper is provided for the sole use of FedCIRC and Global Integrity customers. This version of the document is "Interim Final"; the information contained herein is accurate to the best of our knowledge at the time of this writing. Global Integrity will continue to research this subject and will provide additional information, as it becomes available. This is primarily an idea document, and is not meant to provide exact solutions applicable to every environment. Administrators should apply normal site procedures for the review, testing, and implementation of hardware and software changes. Site specific conditions will impact the design and effectiveness of various measures. Global Integrity makes no warranty, express or implied, as to the effectiveness or impact of the items discussed herein. Global Integrity assumes no liability for the use of information in this paper. Use at your own risk.

This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

Executive Summary

This paper addresses potential defensive measures to take against Distributed Denial of Service attacks; recommendations are applicable both prior to and during an active attack. Recommendations operate at different levels and in most cases multiple recommendations can and should be applied. These techniques should be fully researched in the specific context in which they will be used.

1. Background

Distributed Denial of Service tools utilize compromised systems to launch coordinated DoS attacks against a limited number of targets. The combined bandwidth of multiple attacking agents creates a bandwidth intensive attack and presents new defensive challenges.

The first defensive step is to employ available protection against the known Denial of Service attacks. In addition, other defenses are presented here which can be combined to establish a broad and flexible defensive posture. Note that there is no single defense for DDOS attacks. Depending on the attack and the target architecture, various defenses may be more or less effective. It is critical that:

- Critical systems are documented.
- An attack is quickly detected and identified.
- Incident response procedures are in place.
- Alternative defenses are defined and available.

Note that an attack against one site may adversely affect an upstream provider, and hence affect non-target sites. The recommendations that follow are designed to defend against known attacks as we understand them and may not be effective against future or unpublished attacks.

2. DoS Specific Defenses

Known DDOS tools implement a number of Denial of Service attacks; in most cases, these attacks have been known and widely discussed for a year or more, and vendors have developed defenses against them. The DDOS attacks simply increase the bandwidth of the attacks, initially, potential targets should employ defensive measures to address the specific DoS attacks that a target may experience. The attack tools reviewed were Trin00, TFN, TFN2K, and stacheldraht. Each of these tools implements some combination of SYN flood, UDP flood, ICMP flood, Smurf, and Fraggle attacks (discussed below). Some tools may launch an attack which is a random combination of several different DoS attacks.

2.1. SYN Flood

2.1.1. Attack Description

TCP packets with the SYN bit set are sent to the target; the IP source address is spoofed. The target responds with a SYN-ACK packet, but the source never replies (assuming the spoofed source IP address system is unreachable). These half-open connections quickly exhaust resources on the server (as few as several packets per minute can effect

this attack under default configurations). Recent attacks have used a large number of SYN flood packets, making this a bandwidth intensive attack as well. We have seen SYN floods directed at single ports and at a range of ports.

2.1.2. Defenses

2.1.2.1. Cisco Routers

Cisco routers have a feature called "TCP Intercept" designed specifically to combat SYN flood attacks. To summarize, the TCP Intercept software intercepts TCP SYN packets and establishes a connection with the client on behalf of the destination server, and if successful, establishes the connection with the server on behalf of the client and knits the two half-connections together transparently. Thus, connection attempts from unreachable hosts (i.e., SYN flood traffic) will never reach the server. The router is designed to handle a much larger number of potential connections than the server, but note that the router will consume more resources when using TCP Intercept. For more information and configuration details, see:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113ed_cr/secur_c/scprt3/scdenial.htm

Cisco routers also have a feature called "Committed Access Rate (CAR)". The router will limit the bandwidth consumed by certain types of traffic (configurable via an extended access control list). This can be used to limit the bandwidth consumed by SYN packets, so that non-SYN packets (i.e., legitimate established connections) will have bandwidth available. The downside to this approach is that it will be difficult for a legitimate client to establish a new connection while the target is under attack. One technique using CAR is to permit unrestricted access to a specific set of known critical clients and apply CAR to others.

2.1.2.2. CheckPoint FW-1

FW-1 has a feature called "SYN Defender" designed specifically to combat SYN flood attacks. The concept is the same as Cisco's TCP Intercept: SYN packets are intercepted and validated on behalf of the server; attack packets are denied at the FW-1 or attack-caused half-open server connections are aggressively reset on the server to quickly free resources. As with Cisco's TCP Intercept, the FW-1 will consume additional resources when SYN

This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

Defender is enabled. For more information and configuration details, see:

<http://www.checkpoint.com/products/firewall-1/syndefender.html>

2.1.2.3. Other Firewalls and Routers

Other firewalls and routers may have similar defensive mechanisms; consult your vendor or documentation for details.

2.2. ICMP Flood

2.2.1. Attack Description

Large numbers of ICMP packets (usually echo request) are sent to the target, overwhelming available bandwidth and/or system resources.

2.2.2. Defenses

ICMP Echo Request packets are not required in most architectures. Routers and other filtering devices can drop ICMP Echo Request packets to prevent the packets from reaching the intended target. Other ICMP protocol packets may be more important in some architectures; we do not recommend filtering all ICMP packets unless necessary to defend against an attack.

2.3. UDP Flood

2.3.1. Attack Description

Large numbers of UDP packets are sent to the target, overwhelming available bandwidth and/or system resources. The UDP flood code is the only DoS attack included in the original Trin00 code, and sends the stream of packets to random UDP ports on the target. This attack could be modified to target a single UDP port, such as 53 (DNS).

2.3.2. Defenses

Routers and other filtering devices can drop packets to non-required UDP ports. This will be effective if the attack targets a wide range of UDP ports, but not if the attack targets required UDP ports (like DNS). In the case of required traffic, CAR (discussed above in 2.1.2.1) can be used to protect the target, although legitimate traffic restricted by CAR may have difficulty getting through. This approach can be useful if,

This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

for example, a DNS server also provides other services: the attack traffic to UDP port 53 will not prevent other legitimate traffic from reaching the target. In the case of DNS, it may also be possible to deny all traffic to UDP port 53 and rely on DNS secondary servers to provide DNS services.

2.4. Smurf

2.4.1. Attack Description

This is a variant of the ICMP flood described above. The attacker sends ICMP Echo Request packets to the broadcast address of a network with a spoofed IP address of the target. If not properly configured to prevent this, the target of the initial broadcast packets (called an amplifier) will distribute the ICMP Echo Request to all hosts on the network; each of these hosts will respond with an Echo Reply packet to the end target. This generates a large number of ICMP Echo Reply packets, overwhelming available bandwidth and/or system resources.

2.4.2. Defenses

Similar to ICMP flood attacks, routers and other filtering devices can drop ICMP Echo Reply packets without adverse effects. In the case of a Smurf attack, it is also possible to drop only "naked" ICMP Echo Reply packets; i.e., ICMP Echo Reply packets without corresponding ICMP Echo Request packets. Note that this will require more resources on the filtering system; the functionality provided may not offset the performance impact. The attack traffic which reaches the end target will not be spoofed, so it is also possible to filter on source IP addresses.

To avoid being used as an amplifier, system administrators should disable a system's ability to respond to these directed broadcast packets. For configuration information for various systems, see:

<http://users.quadrunner.com/chuegen/smurf.cgi>

2.5. Fraggle

2.5.1. Attack Description

This is a variant of the Smurf attack described above. The attacker sends UDP Echo packets (UDP port 7) to the broadcast address of a network with a spoofed IP address of the target. If not properly

configured to prevent this, the target of the initial broadcast packets (called an amplifier) will distribute the UDP Echo packets to all hosts on the network; each of these hosts will respond with a UDP Echo packet to the end target. This generates a large number of UDP Echo packets, overwhelming available bandwidth and/or system resources. Other UDP and TCP services such as Chargen can be used in a similar manner. A simpler version of this attack simply sends traffic directly to available UDP or TCP services to occupy available bandwidth and resources.

2.5.2. Defenses

Similar to Smurf attacks, routers and other filtering devices can drop UDP/TCP Echo (and Chargen and Discard) packets without adverse effects. As with Smurf attacks, the attack traffic reaching the end target will not be spoofed; filtering may be applied based on source IP addresses (although there is probably no reason to accept Echo, Chargen, or Discard packets from any source).

To avoid being used as an amplifier, system administrators should disable a system's ability to respond to directed broadcast packets. For configuration information for various systems, see:

<http://users.quadrunner.com/chuegen/smurf.cgi>

Also, system administrators should disable these UDP and TCP services (Echo, Chargen, and Discard) unless necessary (in most circumstances, they are not necessary). Consult your system documentation for configuration details. For Cisco routers, the configuration command is:

```
no service udp-small-servers  
no service tcp-small-servers
```

2.6. Other DoS Attacks

Attackers may add additional known (and possibly unknown) DoS attacks to existing DDOS tools. These cases will require an identification of the attack and application of appropriate specific and general defenses.

3. Upstream Filters

In many cases, the attack will overwhelm available bandwidth. When the attack traffic can be identified and filtered (such as a SYN flood against a wide range of

This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

ports), it will be beneficial to apply filtering as far upstream as possible. This may require the cooperation of your service provider and is architecture dependent.

4. Dedicated Defensive Hardware

Depending on your architecture and the particular type of attack, it may be beneficial to install dedicated hardware to defend against the attack as far upstream as possible. For example, your network provider may allow you to install a powerful router in your network path where your bandwidth is high (prior to leaving your network provider's architecture). This router can normally operate in pass-through mode, but can be reconfigured in the event of an attack. For example, such a dedicated router could perform filtering or SYN flood defense prior to the traffic reaching your lower bandwidth connections.

5. IP Hopping

The known attack tools target a specific IP address or addresses and do not perform DNS lookups to resolve an IP address from a name. Therefore, it may be possible to avert an attack by changing the IP address used to access the target. The attacker may not monitor the attack, in which case a single IP address change will thwart the attack. If the attacker does monitor the attack, then repeated and automated IP address changes may be performed to avert or mitigate the effects of the attack. Two approaches can be taken: (1) have the mechanisms for IP address changing in place and only enable them if attacked, or (2) conduct regular and frequent IP address changes as a matter of course, even if an attack is not present. Option 1 requires a detection capability and may incur a lag time during which the attack may be effective, but this defensive capability will not be exposed to the attacker until needed. Option 2 may be able to thwart the attack entirely, as the attacker may never be able to keep up (including the initial attack), but this defensive capability will be exposed to an astute attacker prior to the attack. A general procedure with variations is presented below; this is designed to be quick, possibly automated, require a minimum of reconfiguration, and impact legitimate clients as little as possible.

5.1. DNS Changes

Modify DNS entries to change the IP address for critical servers; do not reconfigure the servers themselves. The TTL for DNS records should be set to a small value (several minutes, if possible) **in advance**. Based on the additional amount of DNS traffic that this may generate, consider moving the DNS servers to a dedicated link separate from the potential victim. Also ensure that DNS secondary servers are available and distributed (completely different network paths).

5.2. Network Address Translation

To avoid server reconfiguration, the new IP addresses must be translated to the existing server IP addresses. If Network Address Translation (NAT) is in place, it can simply be reconfigured. If NAT is not in place, then install that capability (a dedicated router or routers, for example). Alternatively, environments using load balancing can reconfigure the load balancer to accommodate the new IP addresses.

5.3. Filter the Old IP Address

Traffic to the old IP address will only be attack traffic and clients using stale DNS entries. This traffic can be filtered as far upstream as possible with minimal effects on legitimate clients. A short DNS entry TTL will minimize the number of clients using stale DNS entries. If IP address hopping is used when not under attack, the effects of lagging DNS entry propagation can be mitigated by overlapping the permitted traffic filters (i.e., permit traffic to stale DNS entries for some fixed amount of time). One alternative or enhancement to filtering which may push the attack traffic stopping point further upstream is to stop advertising routes for the old IP address. As the lack of a route for this traffic is propagated, the attack traffic will be rejected further from the target (further than possible with filtering). This approach is dependent on IP addressing and routing schemes.

5.4. Different IP Address Blocks and Different Links

The most aggressive defensive approach would use different IP address blocks and different physical links when changing IP addresses. If the attacker stops one attack when reconfiguring (instead of just adding targets), then the IP address changes can rotate between two different setups. If the attacker simply adds IP targets as the IP addresses are changed, more IP address and/or link possibilities are required, but the effect of the attack can be diluted since non-current IP addresses can be filtered and/or routes disabled.

5.5. DNS Server Defense

The IP address hopping defense relies on DNS to propagate IP address changes; the attacker may target the DNS server itself to counter the defense. The most likely attack against the DNS server would be a UDP port 53 flood or a TCP SYN port 53 attack. Since the attack traffic could be made nearly indistinguishable from legitimate traffic, filtering defenses may not be effective. Several other defenses are available:

- Put the primary DNS server on a dedicated link (if the bandwidth of the attack is fairly low).
- Establish backup primary DNS servers at alternate locations.

This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

- Establish multiple secondary DNS servers with independent network paths.
- Use an invisible primary DNS server (unadvertised, possibly not on the network at all) with dedicated links to the secondary DNS servers (dialup, for example).
- Establish non-advertised secondary DNS servers that can be advertised if necessary.

5.6. The Persistent Attacker

A persistent attacker who follows the IP address changes presents the opportunity for some defensive measures. For example (see [1]):

- More control traffic creates a greater opportunity for traceback.
- Monitoring traffic creates an opportunity for traceback.
- If the attacker's monitoring mechanism can be identified, false information may be provided.
- The attacker must expend finite resources to keep the attack persistent, which may not be worth it for him/her/them.

6. URL Redirect

If the attack is against a web server, a redirection defense may work. See Dr. Fred Cohen's paper on the subject [1]; a summary follows. Change the DNS entry for the web server to point to a dedicated system on a high bandwidth connection independent of the real web server. This dedicated system provides only a URL redirect to the real web server; legitimate connections will follow the redirect, and other traffic (most notably SYN flood traffic) will not. The attacker can easily determine the IP address of the real web server; to defend, the redirect page should change the IP address pointed to in coordination with NAT in front of the real web server. This technique is similar to the IP address hopping technique described above, but has the advantage that DNS records are not used to propagate IP address changes, thus avoiding legitimate clients using stale DNS entries. To further frustrate the attacker, the DNS entry for the dedicated redirect server could be changed at regular intervals to avoid successful bandwidth consumption attacks against the dedicated redirect server.

7. Alternate Limited Access

In certain situations, access for specific clients may be more important than general access. In such cases, separate dedicated access can be set up for those clients independent of general access (or non-critical clients can be filtered). Noted in [1].

8. General Recommendations

This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

The following recommendations are generally applicable:

- Arrange for excess bandwidth, whether in hot standby mode, during normal operations, or provided as requested from the site's network service provider(s).
- Build redundancy into the potential target architecture, both in terms of systems and networks.
- Separate traffic wherever possible; for example, use one provider for external access to site web servers and another provider for internal network access to the Internet.
- Contact the site network service provider(s) regarding their:
 - implemented DoS defenses
 - available DoS defenses
 - emergency contact information
 - upstream providers
- Filter known bad traffic, inbound and outbound. For example:
 - RFC 1918 Addresses (Private Address Space). Under most circumstances, this traffic will never arrive from external sources:
 - 10.0.0.0/8
 - 172.16.0.0/12
 - 192.168.0.0/16
 - Broadcast traffic. Traffic with source or destination broadcast addresses (fourth octet 0 or 255) is probably not legitimate or required traffic.
 - Loopback. This traffic (127.0.0.0/8) should not appear on the network.
 - For egress filtering, spoofed source address traffic should be dropped (and probably logged)..

9. References and Additional Information

- [1] Cohen, Fred. *Managing Network Security: Countering DCAs*.
<http://all.net/journal/netsec/0004.html>. Date Unknown.
- [2] Cisco. *Configuring TCP Intercept (Prevent Denial-of-Service Attacks)*.
http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113ed_cr/secur_c/scprt3/scdenial.htm. Date Unknown.
- [3] Huegen, Craig. *The Latest in Denial of Service Attacks: "Smurfing Description and Information to Minimize Effects*. <http://users.quadrunner.com/chuegen/smurf.cgi>,
chuegen@quadrunner.com. February 8, 2000.
- [4] CheckPoint. *TCP SYN Flooding Attack and the FireWall-1 SYNDefender*.
<http://www.checkpoint.com/products/firewall-1/syndefender.html>. October 1996.
- [5] Dittrich, Dave. *Distributed Denial of Service Tool Analyses*.
<http://www.washington.edu/People/dad/>. 1999-2000.

This is a work-in-progress and presents techniques which have not been fully tested. Follow established site procedures for review, testing, and implementation of hardware and software changes.

10.Point of Contact

Jim Jones
Senior Information Security Analyst
Global Integrity Corporation, REACT Services
Mailstop 2-3
12100 Sunset Hills Road
Reston, VA 20190

jim.jones@globalintegrity.com
(703) 375-2487